The Capacity of Robust Private Information Retrieval With Colluding Databases

Hua Sun, *Member, IEEE*, and Syed Ali Jafar, *Fellow, IEEE*

***Abstract*— Private information retrieval (PIR) is the problem of retrieving as efficiently as possible, one out of K messages from N non-communicating replicated databases (each holds all K messages) while keeping the identity of the desired message index a secret from each individual database. The information theoretic capacity of PIR (equivalently, the reciprocal of minimum download cost) is the maximum number of bits of desired information that can be privately retrieved per bit of downloaded information. T-private PIR is a generalization of PIR to include the requirement that even if any T of the N databases collude, the identity of the retrieved message remains completely unknown to them. Robust PIR is another generalization that refers to the scenario where we have $M \geq N$ databases, out of which any $M - N$ may fail to respond. For K messages and $M \geq N$ databases out of which at least some N must respond, we show that the capacity of T-private and Robust PIR is $(1 + T/N + T^2/N^2 + \cdots + T^{K-1}/N^{K-1})^{-1}$. The result includes as special cases the capacity of PIR without robustness ($M = N$) or T-privacy constraints ($T = 1$).**

***Index Terms*— Capacity, private information retrieval, colluding databases, unresponsive databases.**

I. Introduction

THE private information retrieval (PIR) problem is motivated by the desire to protect the privacy of a user against data providers. Besides its direct applications in data privacy, it is intimately related to many fundamental problems in cryptography, e.g., oblivious transfer [1], instance hiding [2]–[4], secure multiparty computation [5], and secret sharing schemes [6], [7]. The significance of PIR also extends beyond security, through its fundamental connections to other prominent topics such as locally decodable codes [8] and batch codes [9] in coding theory, relationships between communication and computation [10] in complexity theory, and most recently blind interference alignment [11] in wireless communications. In fact most constructions of locally decodable codes are translated directly from PIR schemes. Through the connections between locally decodable and locally recoverable codes [12], PIR also connects to distributed data storage repair [13] and index coding [14], which in turn encompass all of network coding [15]. Therefore PIR represents an important focal point to tackle significant challenges across these fields.

The goal of PIR is to find the most efficient way for a user to retrieve a desired message from a set of N distributed databases, each of which stores all K messages, without revealing anything (in the information theoretic sense)[1] about which message is being retrieved, to any individual database. The PIR problem was initially studied in the setting where each message is one bit long [8], [19]–[23], where the cost of a PIR scheme is measured by the total amount of communication between the user and the databases, i.e., the sum of communications from the user to the databases (upload) and from the databases to the user (download). What is pursued in this work is the traditional Shannon theoretic formulation, where message size is allowed to be arbitrarily large, and therefore the upload cost is negligible compared to the download cost [20], [24]. The information theoretic capacity of PIR is the maximum number of bits of desired information that can be privately retrieved per bit of downloaded information. Equivalently, it is the reciprocal of the minimum possible download cost per bit of desired message. In [25], we showed that the information theoretic capacity of PIR, for arbitrary number of messages K and arbitrary number of databases N is $\left(1 + 1/N + 1/N^2 + \cdots + 1/N^{K-1}\right)^{-1}$.

There are several interesting extensions of PIR that explore its limitations under additional constraints. These include extensions where up to T of the N databases may collude [26], [27] (T-private PIR); where some of the databases may not respond [28] (Robust PIR); where both the privacy of the user and the databases must be protected [1] (Symmetric PIR); where only one database holds all the messages and all other databases hold independent information [29]; where retrieval operations are unsynchronized [30]; and where beyond communications, computation is also a concern [31]. There is also much recent work in the distributed storage setting [24], [32]–[34] (the databases form a distributed storage system) where the main focus is on how the coding of the storage system works jointly with PIR.

In this work, we mainly consider T-private PIR in the Shannon theoretic setting, where we have an arbitrary number

Manuscript received November 8, 2016; revised October 31, 2017; accepted November 6, 2017. Date of publication November 24, 2017; date of current version March 15, 2018. This work was supported in part by NSF under Grant CCF-1617504, Grant CCF-1317351, and Grant CNS-1731384, in part by ONR under Grant N00014-16-1-2629, and in part by ARO under Grant W911NF-16-1-0215. This paper was presented in part at the IEEE GlobalSIP 2016.

H. Sun is with the Department of Electrical Engineering, University of North Texas, Denton, TX 76203, USA (e-mail: hua.sun@unt.edu).

S. A. Jafar is with the Center for Pervasive Communications and Computing, Department of Electrical Engineering and Computer Science, University of California at Irvine, Irvine, CA 92697, USA (e-mail: syed@uci.edu).

Communicated by A. Khisti, Associate Editor for Shannon Theory.

Digital Object Identifier 10.1109/TIT.2017.2777490

[1] There is another line of research, where privacy needs to be satisfied only for computationally bounded databases [16]–[18].

of messages (K), arbitrary number of databases (N), each database stores all the messages, the messages are allowed to be arbitrarily large, and the privacy of the desired message index must be guaranteed even if any T of the N databases collude. The main contribution of this work is to show that the information theoretic capacity of T-private PIR is $\left(1+T/N+T^2/N^2+\cdots+T^{K-1}/N^{K-1}\right)^{-1}$.

We further consider the extension to *robust* T-private PIR, where we have $M \geq N$ databases, out of which any $M-N$ databases may not respond, so that with answers from any N databases, we need to ensure both privacy and correctness. In this context, the contribution of this work is to show that the information theoretic capacity of robust T-private PIR remains the same as that of T-private PIR, i.e., there is no capacity cost from not knowing in advance *which* N databases will respond.

Notation: For $n_1, n_2 \in \mathbb{Z}, n_1 \leq n_2$, define the notation $[n_1 : n_2]$ as the set $\{n_1, n_1+1, \cdots, n_2\}$ and $(n_1 : n_2)$ as the vector $(n_1, n_1+1, \cdots, n_2)$. For an index set $\mathcal{I} = \{i_1, i_2, \cdots, i_n\}$, the notation $A_{\mathcal{I}}$ represents the set $\{A_i : i \in \mathcal{I}\}$. For an index vector $\mathcal{I} = (i_1, i_2, \cdots, i_n)$, the notation $A_{\mathcal{I}}$ represents the vector $(A_{i_1}, A_{i_2}, \cdots, A_{i_n})$. For a matrix S, the notation $S[\mathcal{I}, :]$ represents the submatrix of S formed by retaining only the rows corresponding to the elements of the vector $\mathcal{I}$. The notation $X \sim Y$ is used to indicate that X and Y are identically distributed.

II. Problem Statement

A. T-Private PIR

Consider K independent messages $W_1, \cdots, W_K$ of size L bits each.

$$H(W_1, \cdots, W_K) = H(W_1) + \cdots + H(W_K), \tag{1}$$

$$H(W_1) = \cdots = H(W_K) = L. \tag{2}$$

There are N databases. Each database stores all the messages $W_1, \cdots, W_K$. A user wants to retrieve $W_k, k \in [1 : K]$ subject to T-privacy, i.e., without revealing anything about the message identity, k, to any colluding susbset of up to T out of the N databases.

To retrieve W_k privately, the user generates N queries $Q_1^{[k]}, \cdots, Q_N^{[k]}$, where the superscript denotes the desired message index. Since the queries are generated with no knowledge of the realizations of the messages, the queries must be independent of the messages,

$$I(W_1, \cdots, W_K; Q_1^{[k]}, \cdots, Q_N^{[k]}) = 0. \tag{3}$$

The user sends query $Q_n^{[k]}$ to the n-th database, $\forall n \in [1 : N]$. Upon receiving $Q_n^{[k]}$, the n-th database generates an answering string $A_n^{[k]}$, which is a deterministic function of $Q_n^{[k]}$ and the data stored (i.e., all messages $W_1, \cdots, W_K$),

$$H(A_n^{[k]} | Q_n^{[k]}, W_1, \cdots, W_K) = 0. \tag{4}$$

Each database returns to the user its answer $A_n^{[k]}$. From all answers $A_1^{[k]}, \cdots, A_N^{[k]}$, the user can decode the desired message W_k,

$$\text{[Correctness]}\ H(W_k | A_1^{[k]}, \cdots, A_N^{[k]}, Q_1^{[k]}, \cdots, Q_N^{[k]}) = 0. \tag{5}$$

To satisfy the privacy constraint that any T colluding databases learn nothing about the desired message index k information theoretically, information available to any T databases (queries, answers and the stored messages) must be independent of k. Let $\mathcal{T}$ be a subset of $[1 : N]$ and its cardinality be denoted by $|\mathcal{T}|$. $Q_{\mathcal{T}}^{[k]}$ represents the subset $\{Q_n^{[k]}, n \in \mathcal{T}\}$. $A_{\mathcal{T}}^{[k]}$ is defined similarly. To satisfy the T-privacy requirement we must have

$$\text{[Privacy]}\ \ I(Q_{\mathcal{T}}^{[k]}, A_{\mathcal{T}}^{[k]}, W_1, \cdots, W_K; k) = 0, \quad \forall \mathcal{T} \subset [1 : N], |\mathcal{T}| = T. \tag{6}$$

To underscore that any set of T or fewer answering strings is independent of the desired message index, we may suppress the superscript and write $A_{\mathcal{T}}$ directly instead of $A_{\mathcal{T}}^{[k]}$, and express the elements of such a set as A_n instead of $A_n^{[k]}$.

The metric that we study in this paper is the PIR rate,[2] which characterizes how many bits of desired information are retrieved per downloaded bit. Note that the PIR rate is the reciprocal of download cost. The rate R of a PIR scheme is defined as follows.

$$R \triangleq \frac{L}{D} \tag{7}$$

where D is the expected value of the total number of bits downloaded by the user from all the databases. The capacity, C, is the supremum of R over all PIR schemes.

B. Robust T-Private PIR

The robust T-private PIR problem is defined similar to the T-private PIR problem. The only difference is that instead of N databases, we have $M \geq N$ databases, and the correctness condition needs to be satisfied when the user collects *any* N out of the M answering strings.

III. Main Result: Capacity of Robust T-Private PIR

The following theorem states the main result.

Theorem 1: For T-private PIR with K messages and N databases, the capacity is

$$C = \left(1+T/N+T^2/N^2+\cdots+T^{K-1}/N^{K-1}\right)^{-1}. \tag{8}$$

The capacity of PIR with T colluding databases generalizes the case without T-privacy constraints, where $T=1$ [25]. The capacity is a strictly decreasing function of T. When $T=N$, the capacity is $1/K$, meaning that the user has to download all K messages to be private, as in this case, the colluding databases are as strong as the user. Similar to the $T=1$ case, the capacity is strictly decreasing in the number of messages, K, and strictly increasing in the number of databases, N. When the number of messages approaches infinity, the capacity approaches $1-T/N$, and when the number of databases approaches infinity (T remains

[2] In the Shannon theoretic formulation where the message size is allowed to grow, the upload cost (the length of the query strings) is negligible relative to download cost because it does not scale with message size. For example, if the message size is doubled (e.g., double the value of L in the schemes presented in this paper), the same query applies to both parts of the message. A more detailed treatment may be found in Proposition 4.1.1 of [20].

constant), the capacity approaches 1. Finally, note that since the download cost is the reciprocal of the rate, the capacity characterization in Theorem 1 equivalently characterizes the optimal download cost per message bit for T-private PIR as $\left(1+T/N+T^2/N^2+\cdots+T^{K-1}/N^{K-1}\right)$ bits. Note that when $N \neq T$, the capacity expression can be equivalently expressed as $(1-\frac{T}{N})/\left(1-\left(\frac{T}{N}\right)^K\right)$.

The capacity-achieving scheme that we construct for T-private PIR, generalizes easily to incorporate robustness constraints. As a consequence, we are also able to characterize the capacity of robust T-private PIR. This result is stated in the following theorem.

Theorem 2: The capacity of robust T-private PIR is

$$C = \left(1+T/N+T^2/N^2+\cdots+T^{K-1}/N^{K-1}\right)^{-1}. \quad (9)$$

Since the capacity expressions are the same, we note that there is no capacity penalty from not knowing in advance which N databases will respond. Even though this uncertainty increases as M increases, capacity is not a function of M. However, we note that the communication complexity of our capacity achieving scheme does increase with M.

Remark: The capacity results in both Theorem 1 and Theorem 2 extend to the ϵ-error case. Please refer to the Appendix for details.

IV. Proof of Theorem 1 and Theorem 2: Achievability

The achievability of the two theorems follows along similar lines, so we present the proofs together in this section.

There are two key aspects of the achievable scheme – 1) the query structure, and 2) the specialization of the query structure to ensure T-privacy and correctness. While the query structure is different from the $T=1$ setting of [25], it draws upon the iterative application of the same three principles that were identified in [25]. These principles are listed below.

(1) *Enforcing Symmetry Across Databases*
(2) *Enforcing Message Symmetry within the Query to Each Database*
(3) *Exploiting Previously Acquired Side Information of Undesired Messages to Retrieve New Desired Information*

The specialization of the structure to ensure T-privacy and correctness is another novel element of the achievable scheme. To illustrate how these ideas work together in an iterative fashion, we will present a few simple examples corresponding to small values of K, M, N and T, and then generalize it to arbitrary K, M, N and T. Let us begin with a lemma.

Lemma 1: Let $S_1, S_2, \cdots, S_K \in \mathcal{F}_q^{\alpha\times\alpha}$ be K random matrices, drawn independently and uniformly from all $\alpha\times\alpha$ full-rank matrices over $\mathcal{F}_q$. Let $G_1, G_2, \cdots, G_K \in \mathcal{F}_q^{\beta\times\beta}$ be K invertible square matrices of dimension $\beta\times\beta$ over $\mathcal{F}_q$. Let $\mathcal{I}_1, \mathcal{I}_2, \cdots, \mathcal{I}_K \in \mathbb{N}^{\beta\times 1}$ be K index vectors, each containing β distinct indices from $[1:\alpha]$. Then

$$(G_1S_1[\mathcal{I}_1,:], G_2S_2[\mathcal{I}_2,:], \cdots, G_KS_K[\mathcal{I}_K,:]) \sim (S_1[(1:\beta),:], S_2[(1:\beta),:], \cdots, S_K[(1:\beta),:]) \quad (10)$$

where $S_i[\mathcal{I}_i,:], i\in[1:K]$ are $\beta\times\alpha$ matrices comprised of the rows of S_i with indices in $\mathcal{I}_i$.

Proof: We wish to prove that the left hand side of (10) is identically distributed (recall that the notation $X\sim Y$ means that X and Y are identically distributed) to the right hand side of (10). Because the rank of a matrix does not depend on the ordering of the rows, we have

$$(S_1[\mathcal{I}_1,:], S_2[\mathcal{I}_2,:], \cdots, S_K[\mathcal{I}_K,:]) \sim (S_1[(1:\beta),:], S_2[(1:\beta),:], \cdots, S_K[(1:\beta),:]).$$

Since S_i are picked uniformly from all full-rank matrices, conditioned on any feasible value of the remaining rows $S_i[(\beta+1:\alpha),:]$, the first β rows $S_i[(1:\beta),:]$ are uniformly distributed over all possibilities that preserve full-rank for S_i. Now note that the mapping from $S_i[(1:\beta),:]$ to $G_iS_i[(1:\beta),:]$ is bijective, and $S_i[(1:\beta),:]$ spans the same row space as $G_iS_i[(1:\beta),:]$, i.e., replacing $S_i[(1:\beta),:]$ with $G_iS_i[(1:\beta),:]$, preserves S_i as a full-rank matrix. Therefore, conditioned on any feasible $S_i[(\beta+1:\alpha),:]$, the set of feasible values of $S_i[(1:\beta),:]$ is the same as the set of feasible $G_iS_i[(1:\beta),:]$ values. Therefore, $G_iS_i[(1:\beta),:]$ is also uniformly distributed over the same set. Finally, since the S_i are chosen independently, the statement of Lemma 1 follows. ∎

In the following, we present 3 examples. In the first two examples, all databases respond, while the last example is on the robust setting.

A. $K=2$ Messages, $M=N=3$ Databases, $T=2$ Colluding Databases

The capacity for this setting, is $C=\left(1+\frac{2}{3}\right)^{-1}=\frac{3}{5}$.

1) Query Structure: We begin by constructing a query structure, which will then be specialized to achieve correctness and privacy. Without loss of generality, let $[a_k]$ denote the symbols of the desired message, and $[b_k]$ the symbols of the undesired message.

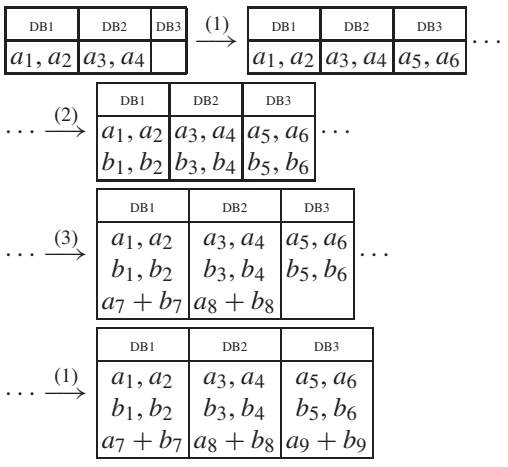

We start by requesting the first $T^{K-1} = 2$ symbols from each of the first $T = 2$ databases: a_1, a_2 from DB1, and a_3, a_4 from DB2. Applying database symmetry, we simultaneously request a_5, a_6 from DB3. Next, we enforce message symmetry, by including queries for $b_1, \cdots, b_6$ as the counterparts for $a_1, \cdots, a_6$. Now consider the first $T = 2$ databases, i.e., DB1 and DB2, which can potentially collude with each other. Unknown to these databases the user has acquired two symbols of external side information, b_5, b_6, comprised of undesired message symbols received from DB3. Splitting the two symbols of external side information among DB1 and DB2 allows the user one symbol of side information for each of DB1 and DB2 that it can exploit to retrieve new desired information symbols. In our construction of the query structure, we will assign new labels (subscripts) to the external side information exploited within each database, e.g., b_7 for DB1 and b_8 for DB2, with the understanding that eventually when the dependencies within the structure are specialized, b_7, b_8 will turn out to be functions of previously acquired side information. Using its assigned side information, each DB acquires a new symbol of desired message, so that DB1 requests $a_7 + b_7$ and DB2 requests $a_8 + b_8$. Finally, enforcing symmetry across databases, DB3 requests $a_9 + b_9$. At this point, the construction is symmetric across databases, the query to any database is symmetric in itself across messages, and the amount of side information exploited within any T colluding databases equals the amount of side information available external to those T databases. So the skeleton of the query structure is complete.

Note that if DB1 and DB2 collude, then the external side information is b_5, b_6, so we would like the side information that is exploited by DB1 and DB2, i.e., b_7, b_8 to be functions of the external side information that is available, i.e., b_5, b_6. However, since *any* $T = 2$ databases can collude, it is also possible that DB1 and DB3 collude instead, in which case we would like b_7, b_9 to be functions of side information that is external to DB1 and DB3, i.e., b_3, b_4. Similarly, if DB2 and DB3 collude, then we would like b_8, b_9 to be functions of b_1, b_2. How to achieve such dependencies in a manner that preserves privacy and ensures correctness is the remaining challenge. Intuitively, the key is to make b_7, b_8, b_9 depend on *all* side information $b_1, b_2, \cdots, b_6$ in a generic sense. In other words, we will achieve the desired functional dependencies by viewing $b_1, b_2, \cdots, b_9$ as the outputs of a $(9, 6)$ MDS code, so that any 3 of these b_k are functions of the remaining 6. The details of this specialization are described next.

2) Specialization to Ensure Correctness and Privacy: Let each message consist of $N^K = 9$ symbols from a sufficiently large[3] finite field $\mathbb{F}_q$ (i.e., $L = 9$). The messages W_1, $W_2 \in \mathbb{F}_q^{9\times 1}$ are then represented as 9×1 vectors over $\mathbb{F}_q$. Let $S_1, S_2 \in \mathbb{F}_q^{9\times 9}$ represent random matrices chosen privately by the user, independently and uniformly from all 9×9 full-rank matrices over $\mathbb{F}_q$. Without loss of generality, let us assume that W_1 is the desired message. Define the 9×1 vectors $a_{[1:9]} \in \mathbb{F}_q^{9\times 1}$ and $b_{[1:9]} \in \mathbb{F}_q^{9\times 1}$, as follows

$$a_{[1:9]} = S_1 W_1 \tag{11}$$

$$b_{[1:9]} = \text{MDS}_{9\times 6} S_2[(1:6), :] W_2 \tag{12}$$

where $S_2[(1:6), :]$ is a 6×9 matrix comprised of the first 6 rows of S_2. $\text{MDS}_{9\times 6}$ is the generator matrix of a $(9, 6)$ MDS code (e.g., a Reed Solomon code). The generator matrix does not need to be random, i.e., it may be globally known. Note that because of the MDS property, any 6 rows of $\text{MDS}_{9\times 6}$ form a 6×6 invertible matrix. Therefore, from any 6 elements of $b_{[1:9]}$, all 9 elements of $b_{[1:9]}$ can be recovered. For example, from $b_1, b_2, \cdots, b_6$, one can recover b_7, b_8, b_9. The queries from each database are constructed according to the structure described earlier.

DB1	DB2	DB3
a_1, a_2	a_3, a_4	a_5, a_6
b_1, b_2	b_3, b_4	b_5, b_6
$a_7 + b_7$	$a_8 + b_8$	$a_9 + b_9$

(13)

Correctness is easy to see, because the user recovers $b_{[1:6]}$ explicitly, from which it can recover all $b_{[1:9]}$, thereby allowing it to recover all of $a_{[1:9]}$. Let us see why privacy holds. The queries for any $T = 2$ colluding databases are comprised of 6 variables from $a_{[1:9]}$ and 6 variables from $b_{[1:9]}$. Let the indices of these variables be denoted by the 6×1 vectors $\mathcal{I}_a, \mathcal{I}_b \in \mathbb{N}^{6\times 1}$, respectively, so that,

$$\begin{aligned}(a_{\mathcal{I}_a}, b_{\mathcal{I}_b}) &= (S_1[\mathcal{I}_a, :] W_1, \text{MDS}_{9\times 6}[\mathcal{I}_b, :] S_2[(1:6), :] W_2) \\ &\sim (S_1[(1:6), :] W_1, S_2[(1:6), :] W_2)\end{aligned} \tag{14}$$

where (14) follows from Lemma 1 because $\text{MDS}_{9\times 6}[\mathcal{I}_b, :]$ is an invertible 6×6 matrix. Therefore, the random map from W_1 to $a_{\mathcal{I}_a}$ variables is i.i.d. as the random map from W_2 to $b_{\mathcal{I}_b}$, and privacy is guaranteed. Note that since 9 desired symbols are recovered from a total of 15 downloaded symbols, the rate achieved by this scheme is $9/15 = 3/5$, which matches the capacity for this setting. While this specialization suffices for our purpose (it achieves capacity), we note that further simplifications of the scheme are possible, which allow it to operate over smaller fields and with lower upload cost. Such an example is provided in the conclusion section of this paper.

B. $K = 3$ Messages, $M = N = 3$ Databases, $T = 2$ Colluding Databases

The capacity for this setting, is $C = \left(1 + \frac{2}{3} + (\frac{2}{3})^2\right)^{-1} = \frac{9}{19}$.

1) Query Structure: The query structure is constructed as follows.

DB1	DB2	DB3
a_1, a_2, a_3, a_4	a_5, a_6, a_7, a_8	

$\xrightarrow{(1)}$

DB1	DB2	DB3
a_1, a_2, a_3, a_4	a_5, a_6, a_7, a_8	$a_9, a_{10}, a_{11}, a_{12}$

$\xrightarrow{(2)}$

DB1	DB2	DB3
a_1, a_2, a_3, a_4	a_5, a_6, a_7, a_8	$a_9, a_{10}, a_{11}, a_{12}$
b_1, b_2, b_3, b_4	b_5, b_6, b_7, b_8	$b_9, b_{10}, b_{11}, b_{12}$
c_1, c_2, c_3, c_4	c_5, c_6, c_7, c_8	$c_9, c_{10}, c_{11}, c_{12}$

[3]The requirements on the size of the field have to do with the existence of MDS codes that are used in the construction. In this case $q \geq N^K$ is sufficient. We note that the size of the field of operations may be reduced. Such an example is presented in Section VI.

DB1	DB2	DB3
a_1, a_2, a_3, a_4	a_5, a_6, a_7, a_8	$a_9, a_{10}, a_{11}, a_{12}$
b_1, b_2, b_3, b_4	b_5, b_6, b_7, b_8	$b_9, b_{10}, b_{11}, b_{12}$
c_1, c_2, c_3, c_4	c_5, c_6, c_7, c_8	$c_9, c_{10}, c_{11}, c_{12}$
$a_{13} + b_{13}$	$a_{15} + b_{15}$	
$a_{14} + b_{14}$	$a_{16} + b_{16}$	
$a_{17} + c_{13}$	$a_{19} + c_{15}$	
$a_{18} + c_{14}$	$a_{20} + c_{16}$	

(3) $\longrightarrow$

DB1	DB2	DB3
a_1, a_2, a_3, a_4	a_5, a_6, a_7, a_8	$a_9, a_{10}, a_{11}, a_{12}$
b_1, b_2, b_3, b_4	b_5, b_6, b_7, b_8	$b_9, b_{10}, b_{11}, b_{12}$
c_1, c_2, c_3, c_4	c_5, c_6, c_7, c_8	$c_9, c_{10}, c_{11}, c_{12}$
$a_{13} + b_{13}$	$a_{15} + b_{15}$	$a_{21} + b_{17}$
$a_{14} + b_{14}$	$a_{16} + b_{16}$	$a_{22} + b_{18}$
$a_{17} + c_{13}$	$a_{19} + c_{15}$	$a_{23} + c_{17}$
$a_{18} + c_{14}$	$a_{20} + c_{16}$	$a_{24} + c_{18}$

(1) $\longrightarrow$

DB1	DB2	DB3
a_1, a_2, a_3, a_4	a_5, a_6, a_7, a_8	$a_9, a_{10}, a_{11}, a_{12}$
b_1, b_2, b_3, b_4	b_5, b_6, b_7, b_8	$b_9, b_{10}, b_{11}, b_{12}$
c_1, c_2, c_3, c_4	c_5, c_6, c_7, c_8	$c_9, c_{10}, c_{11}, c_{12}$
$a_{13} + b_{13}$	$a_{15} + b_{15}$	$a_{21} + b_{17}$
$a_{14} + b_{14}$	$a_{16} + b_{16}$	$a_{22} + b_{18}$
$a_{17} + c_{13}$	$a_{19} + c_{15}$	$a_{23} + c_{17}$
$a_{18} + c_{14}$	$a_{20} + c_{16}$	$a_{24} + c_{18}$
$b_{19} + c_{19}$	$b_{21} + c_{21}$	$b_{23} + c_{23}$
$b_{20} + c_{20}$	$b_{22} + c_{22}$	$b_{24} + c_{24}$

(2) $\longrightarrow$

DB1	DB2	DB3
a_1, a_2, a_3, a_4	a_5, a_6, a_7, a_8	$a_9, a_{10}, a_{11}, a_{12}$
b_1, b_2, b_3, b_4	b_5, b_6, b_7, b_8	$b_9, b_{10}, b_{11}, b_{12}$
c_1, c_2, c_3, c_4	c_5, c_6, c_7, c_8	$c_9, c_{10}, c_{11}, c_{12}$
$a_{13} + b_{13}$	$a_{15} + b_{15}$	$a_{21} + b_{17}$
$a_{14} + b_{14}$	$a_{16} + b_{16}$	$a_{22} + b_{18}$
$a_{17} + c_{13}$	$a_{19} + c_{15}$	$a_{23} + c_{17}$
$a_{18} + c_{14}$	$a_{20} + c_{16}$	$a_{24} + c_{18}$
$b_{19} + c_{19}$	$b_{21} + c_{21}$	$b_{23} + c_{23}$
$b_{20} + c_{20}$	$b_{22} + c_{22}$	$b_{24} + c_{24}$
$a_{25} + b_{25} + c_{25}$	$a_{26} + b_{26} + c_{26}$	

(3) $\longrightarrow$

DB1	DB2	DB3
a_1, a_2, a_3, a_4	a_5, a_6, a_7, a_8	$a_9, a_{10}, a_{11}, a_{12}$
b_1, b_2, b_3, b_4	b_5, b_6, b_7, b_8	$b_9, b_{10}, b_{11}, b_{12}$
c_1, c_2, c_3, c_4	c_5, c_6, c_7, c_8	$c_9, c_{10}, c_{11}, c_{12}$
$a_{13} + b_{13}$	$a_{15} + b_{15}$	$a_{21} + b_{17}$
$a_{14} + b_{14}$	$a_{16} + b_{16}$	$a_{22} + b_{18}$
$a_{17} + c_{13}$	$a_{19} + c_{15}$	$a_{23} + c_{17}$
$a_{18} + c_{14}$	$a_{20} + c_{16}$	$a_{24} + c_{18}$
$b_{19} + c_{19}$	$b_{21} + c_{21}$	$b_{23} + c_{23}$
$b_{20} + c_{20}$	$b_{22} + c_{22}$	$b_{24} + c_{24}$
$a_{25} + b_{25} + c_{25}$	$a_{26} + b_{26} + c_{26}$	$a_{27} + b_{27} + c_{27}$

(1) $\longrightarrow$

Starting with $T^{K-1} = 4$ symbols each requested from the first $T = 2$ databases, we proceed through iterative steps (1) and (2) to enforce symmetries across databases and messages. In step (3) we consider the first $T = 2$ databases together (DB1 and DB2) and account for the external side information, which in this case contains 4 symbols from $[b_k]$ and 4 symbols from $[c_k]$. Distributed evenly among DB1 and DB2, this allows a budget of 2 symbols of side information from $[b_k]$ and 2 symbols from $[c_k]$ per database to be exploited to recover new symbols of desired information. Proceeding again through steps (1) and (2) to enforce symmetries across databases and messages, we end up with new downloads that contain only undesired information symbols, which can now be used to download new desired information symbols. Once again, we consider DB1 and DB2 together, and account for the new external side information, $b_{23} + c_{23}, b_{24} + c_{24}$. Thus the external side information is comprised of two symbols, each of which is a sum of the form $b_k + c_k$. Dividing the side information evenly among databases DB1 and DB2, each is assigned one side information symbol of the form $b_k + c_k$ with new labels. Thus, $a_{25} + b_{25} + c_{25}$ is added to the query from DB1, and $a_{26} + b_{26} + c_{26}$ is added to the query from DB2. Finally, applying symmetry across databases, we include $a_{27}+b_{27}+c_{27}$ to the query from DB3. At this point, all symmetries are satisfied, all external and exploited side information amounts are balanced, and therefore, the query structure is complete.

2) Specialization: Let each message consist of $N^K = 27$ symbols from a sufficiently large finite field $\mathbb{F}_q$ (i.e., $L = 27$). The messages $W_1, W_2, W_3 \in \mathbb{F}_q^{27\times 1}$ are then represented as 27×1 vectors over $\mathbb{F}_q$. Let $S_1, S_2, S_3 \in \mathbb{F}_q^{27\times 27}$ represent random matrices chosen privately by the user, independently and uniformly from all 27×27 full-rank matrices over $\mathbb{F}_q$. Without loss of generality, let us assume that W_1 is the desired message. Define 27×1 vectors $a_{[1:27]}, b_{[1:27]}, c_{[1:27]} \in \mathbb{F}_q^{27\times 1}$, as follows

$$a_{[1:27]} = S_1 W_1 \tag{15}$$

$$b_{[1:18]} = \text{MDS}_{18\times 12} S_2[(1:12), :] W_2 \tag{16}$$

$$c_{[1:18]} = \text{MDS}_{18\times 12} S_3[(1:12), :] W_3 \tag{17}$$

$$b_{[19:27]} = \text{MDS}_{9\times 6} S_2[(13:18), :] W_2 \tag{18}$$

$$c_{[19:27]} = \text{MDS}_{9\times 6} S_3[(13:18), :] W_3 \tag{19}$$

where $S_2[(1:18), :]$ is a 18×27 matrix comprised of the first 18 rows of S_2. $\text{MDS}_{18\times 12}$ is the generator matrix of a $(18, 12)$ MDS code, and $\text{MDS}_{9\times 6}$ is the generator matrix of a $(9, 6)$ MDS code. In particular, note that the *same* generator matrix is used in (16) and (17). Similarly, the same generator matrix is used in (18) and (19). This is important because it allows us to write

$$b_{[19:27]} + c_{[19:27]} = \text{MDS}_{9\times 6}\left(S_2[(13:18), :] W_2 + S_3[(13:18), :] W_3\right) \tag{20}$$

so that all 9 elements of the vector $b_{[19:27]} + c_{[19:27]}$ can be recovered from any 6 of its elements, e.g., from $b_{[19:24]} + c_{[19:24]}$ one can also recover $b_{25} + c_{25}, b_{26} + c_{26}, b_{27} + c_{27}$. This observation is the key to understanding the role of interference alignment in this construction. The effective number of *resolvable* undesired symbols is minimized due to interference alignment. For example, b_{19} and c_{19} are always aligned together into one symbol $b_{19} + c_{19}$ in all the downloaded equations. The two are unresolvable from each other and act as effectively one undesired symbol in the downloaded equations, thus reducing the effective number of undesired symbols, so that the same number of downloaded equations can be used to retrieve a greater number of desired symbols. Note also that desired symbols are always resolvable.

These values are plugged into the query structure derived previously.

DB1	DB2	DB3
a_1, a_2, a_3, a_4	a_5, a_6, a_7, a_8	$a_9, a_{10}, a_{11}, a_{12}$
b_1, b_2, b_3, b_4	b_5, b_6, b_7, b_8	$b_9, b_{10}, b_{11}, b_{12}$
c_1, c_2, c_3, c_4	c_5, c_6, c_7, c_8	$c_9, c_{10}, c_{11}, c_{12}$
$a_{13} + b_{13}$	$a_{15} + b_{15}$	$a_{21} + b_{17}$
$a_{14} + b_{14}$	$a_{16} + b_{16}$	$a_{22} + b_{18}$
$a_{17} + c_{13}$	$a_{19} + c_{15}$	$a_{23} + c_{17}$
$a_{18} + c_{14}$	$a_{20} + c_{16}$	$a_{24} + c_{18}$
$b_{19} + c_{19}$	$b_{21} + c_{21}$	$b_{23} + c_{23}$
$b_{20} + c_{20}$	$b_{22} + c_{22}$	$b_{24} + c_{24}$
$a_{25} + b_{25} + c_{25}$	$a_{26} + b_{26} + c_{26}$	$a_{27} + b_{27} + c_{27}$

Correctness is straightforward. Let us see why T-privacy holds. The queries for any $T = 2$ colluding databases are comprised of 18 variables from $a_{[1:27]}$, 12 variables from $b_{[1:18]}$, 6 variables from $b_{[19:27]}$, 12 variables from $c_{[1:18]}$ and 6 variables from $c_{[19:27]}$. Let the indices of these variables be denoted by the vectors $\mathcal{I}_a \in \mathbb{N}^{18\times 1}, \mathcal{I}_{b,12} \in \mathbb{N}^{12\times 1}, \mathcal{I}_{b,6} \in \mathbb{N}^{6\times 1}, \mathcal{I}_{c,12} \in \mathbb{N}^{12\times 1}$ and $\mathcal{I}_{c,6} \in \mathbb{N}^{6\times 1}$, respectively, so that,

$$a_{\mathcal{I}_a} = S_1[\mathcal{I}_a, :]W_1 \tag{21}$$

$$b_{\mathcal{I}_{b,12}} = \text{MDS}_{18\times 12}[\mathcal{I}_{b,12}, :]S_2[(1:12), :]W_2 \tag{22}$$

$$b_{\mathcal{I}_{b,6}} = \text{MDS}_{9\times 6}[\mathcal{I}_{b,6}, :]S_2[(13:18), :]W_2 \tag{23}$$

$$c_{\mathcal{I}_{c,12}} = \text{MDS}_{18\times 12}[\mathcal{I}_{c,12}, :]S_3[(1:12), :]W_3 \tag{24}$$

$$c_{\mathcal{I}_{c,6}} = \text{MDS}_{9\times 6}[\mathcal{I}_{c,6}, :]S_3[(13:18), :]W_3 \tag{25}$$

From Lemma 1, we have

$$\begin{aligned}&(a_{\mathcal{I}_a}, (b_{\mathcal{I}_{b,12}};\ \ b_{\mathcal{I}_{b,6}}), (c_{\mathcal{I}_{c,12}};\ c_{\mathcal{I}_{c,6}}))\\ &\quad\sim (S_1[(1:18), :]W_1, S_2[(1:18), :]W_2, S_3[(1:18), :]W_3)\end{aligned}$$

Thus privacy is guaranteed. Finally, note that since 27 desired symbols are recovered from a total of 57 downloaded symbols, the rate achieved by this scheme is $27/57 = 9/19$, which matches the capacity for this setting.

C. $K = 2$ Messages, $M = 3$ Databases, $N = 2$ Responding Databases, $T = 1$ Colluding Database

The capacity for this setting, is $C = \left(1 + \frac{1}{2}\right)^{-1} = \frac{2}{3}$.

1) Query Structure: We first construct the query structure, following the 3 iterative principles previously used for T-private PIR. Without loss of generality, let $[a_k]$ denote the symbols of the desired message, and $[b_k]$ the symbols of the undesired message.

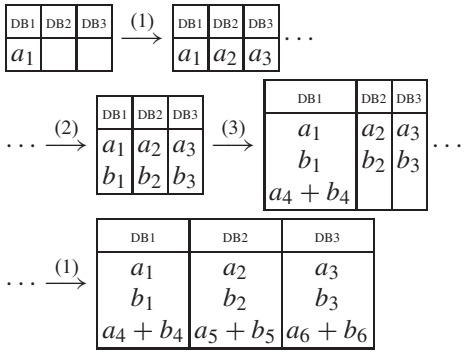

We start by requesting the first $T^{K-1} = 1$ symbol from the first $T = 1$ database, i.e., a_1 from DB1. Applying database symmetry, we simultaneously request a_2 from DB2 and a_3 from DB3. Next, we enforce message symmetry, by including queries for b_1, b_2, b_3 as the counterparts for a_1, a_2, a_3. Note that only $N = 2$ databases may respond. As a result, from the perspective of any individual database, we have only one symbol of external side information (from the other surviving database). We then exploit this side information symbol to retrieve a new desired symbol, i.e., we download $a_4 + b_4$ from DB1, $a_5 + b_5$ from DB2 and $a_6 + b_6$ from DB3. The construction is complete.

We want to ensure that no matter which 2 databases respond, we can gather enough desired symbols to decode the desired message and privacy is preserved to each individual database. These are guaranteed by the following specialization.

2) Specialization to Ensure Correctness and Privacy: Let each message consist of $N^K = 4$ symbols from a sufficiently large field (i.e., $L = 4$). The messages $W_1,\ W_2 \in \mathbb{F}_q^{4\times 1}$ are then represented as 4×1 vectors over $\mathbb{F}_q$. Let $S_1, S_2 \in \mathbb{F}_q^{4\times 4}$ represent random matrices chosen privately by the user, independently and uniformly from all 4×4 full-rank matrices over $\mathbb{F}_q$. Without loss of generality, let us assume that W_1 is the desired message. Define the 6×1 vectors $a_{[1:6]} \in \mathbb{F}_q^{6\times 1}$ and $b_{[1:6]} \in \mathbb{F}_q^{6\times 1}$, as follows

$$a_{[1:6]} = \text{MDS}_{6\times 4}S_1W_1 \tag{26}$$

$$b_{[1:6]} = \text{MDS}_{6\times 2}S_2[(1:2), :]W_2 \tag{27}$$

where $S_2[(1:2), :]$ is a 2×4 matrix comprised of the first 2 rows of S_2. $\text{MDS}_{6\times 4}/\text{MDS}_{6\times 2}$ is the generator matrix of a $(6, 4)/(6, 2)$ MDS code.

DB1	DB2	DB3
a_1	a_2	a_3
b_1	b_2	b_3
$a_4 + b_4$	$a_5 + b_5$	$a_6 + b_6$

(28)

Correctness is easy to see, because after receiving answers from any $N = 2$ databases, the user recovers all $b_{[1:6]}$ (refer to (27)). Then the user subtracts out $b_{[1:6]}$ to recover 4 symbols in $a_{[1:6]}$, from which all $a_{[1:6]}$ are recovered (refer to (26)). The query for any individual database is comprised of 2 variables from $a_{[1:6]}$ and 2 variables from $b_{[1:6]}$. Let the indices of these variables be denoted by the 2×1 vectors $\mathcal{I}_a, \mathcal{I}_b \in \mathbb{N}^{2\times 1}$, respectively, so that,

$$\begin{aligned}&(a_{\mathcal{I}_a}, b_{\mathcal{I}_b})\\ &\quad= (\text{MDS}_{6\times 4}[\mathcal{I}_a, :]S_1W_1, \text{MDS}_{6\times 2}[\mathcal{I}_b, :]S_2[(1:2), :]W_2)\\ &\quad\sim (S_1[(1:2), :]W_1, S_2[(1:2), :]W_2)\end{aligned} \tag{29}$$

where (29) follows from Lemma 1. Therefore, the random map from W_1 to $a_{\mathcal{I}_a}$ variables is i.i.d. as the random map from W_2 to $b_{\mathcal{I}_b}$, and privacy is guaranteed. Note that since 4 desired symbols are recovered from a total of 6 downloaded symbols (from $N = 2$ responding databases), the rate achieved by this scheme is $4/6 = 2/3$, which matches the capacity for this setting.

D. Arbitrary Number of Messages K, Arbitrary Number of Databases M, Arbitrary Number of Responding Databases N, Arbitrary Number of Colluding Databases T

1) Query Structure: For arbitrary K, M, N, T, we follow the same iterative procedure, briefly summarized below.[4]

- Step 1: Initialization. Download T^{K-1} desired symbols each from the first T databases.
- Step 2: Invoke symmetry across databases to determine corresponding downloads from DB $T+1$ to DB M.
- Step 3: Invoke symmetry of messages to determine additional downloaded equations (comprised only of undesired symbols) from each database.
- Step 4: Consider the first T databases together. Divide the new external side information generated in the previous step (note that as $M-N$ databases may not respond, side information is counted from $N-T$ other databases) evenly among the first T databases to determine the side information budget per database. For each side information symbol allocated to a database create an additional query of the same form as the assigned side information (with new labels) combined with a new desired symbol.
- Step 5: Go back to Step 2 and run Step 2 to Step 4 a total of $(K-1)$ times.

2) Specialization: We now map the message symbols to the symbols in the query structure. Let each message consist of N^K symbols from a sufficiently large finite field $\mathbb{F}_q$ (i.e., $L = N^K$). The messages $W_1, \cdots, W_K \in \mathbb{F}_q^{N^K \times 1}$ are represented as $N^K \times 1$ vectors over $\mathbb{F}_q$. Let $S_1, \cdots, S_K \in \mathbb{F}_q^{N^K \times N^K}$ represent random matrices chosen privately by the user, independently and uniformly from all $N^K \times N^K$ full-rank matrices over $\mathbb{F}_q$. Suppose $W_l, l \in [1:K]$, is the desired message.

Consider any undesired message index $k \in [1:K]/\{l\}$, and all distinct $\Delta = 2^{K-2}$ subsets of $[1:K]$ that contain k and do not contain l. Assign distinct labels to each subset, e.g., $\mathcal{K}_1, \mathcal{K}_2, \cdots \mathcal{K}_\Delta$. For each $k \in [1:K]/\{l\}$, define the vector shown at the top of the next page, where $\alpha_i, i \in [1:\Delta]$ is defined as[5] $N(N-T)^{|\mathcal{K}_i|-1}T^{K-|\mathcal{K}_i|}$, each $x^{[k]}_{\mathcal{K}_i}$ is a $\frac{M}{N}\alpha_i \times 1$ vector, and each $x^{[k]}_{\mathcal{K}_i \cup \{l\}}$ is a $\frac{M}{N}(\frac{N-T}{T})\alpha_i \times 1$ vector over $\mathbb{F}_q$.

Now consider the desired message index l, and all distinct $\delta = 2^{K-1}$ subsets of $[1:K]$ that contain l. Assign distinct labels to each subset, e.g., $\mathcal{L}_1, \mathcal{L}_2, \cdots \mathcal{L}_\delta$. Define the vector

$$\begin{bmatrix} x^{[l]}_{\mathcal{L}_1} \\ x^{[l]}_{\mathcal{L}_2} \\ \vdots \\ x^{[l]}_{\mathcal{L}_\delta} \end{bmatrix} = \text{MDS}_{\frac{M}{N}N^K \times N^K} S_l W_l$$

where the length of $x^{[l]}_{\mathcal{L}_i}, i \in [1:\delta]$ is $M(N-T)^{|\mathcal{L}_i|-1}T^{K-|\mathcal{L}_i|}$.

For each non-empty subset $\mathcal{K} \subset [1:K]$ generate the query vector

$$\sum_{k \in \mathcal{K}} x^{[k]}_{\mathcal{K}} \tag{30}$$

Distribute the elements of the query vector evenly among the M databases. This completes the specialized construction of the queries.

The construction has K layers. Over the j-th layer, for each database, there are $(N-T)^{j-1}T^{K-j}\binom{K}{j}$ equations[6] that are comprised of sums of j symbols, out of which $(N-T)^{j-1}T^{K-j}\binom{K-1}{j-1}$ involve desired data symbols.

Suppose the user collects answering strings from any N databases. For each set $\mathcal{K}_i$, from N databases, we download α_i symbols from $x^{[k]}_{\mathcal{K}_i}, k \neq l, i \in [1:\Delta]$, from which we can recover the interference $x^{[k]}_{\mathcal{K}_i \cup \{l\}}$, as they are generated by the generator matrix of a $(\frac{M}{T}\alpha_i, \alpha_i)$ MDS code. After subtracting out all the interference, we are left with N^K desired symbols, from which we can recover the desired message, as the symbols are generated by the generator matrix of a $(\frac{M}{N}N^K, N^K)$ MDS code. Therefore correctness is guaranteed.

Let us see why privacy holds. The queries for any T colluding databases are comprised of TN^{K-1} variables from each $x^{[k]}, k \in [1:K]$. When $k = l$, the TN^{K-1} desired symbols are generated by the generator matrix of a $(\frac{M}{N}N^K, N^K)$ MDS code such that these symbols have full rank. For each $k \neq l$, the TN^{K-1} variables from $x^{[k]}$ consist of α_i variables out of $\frac{M}{T}\alpha_i$ variables $x^{[k]}_{\mathcal{K}_i}, x^{[k]}_{\mathcal{K}_i \cup \{l\}}$, for each set $\mathcal{K}_i, i \in [1:\Delta]$. Note that these α_i variables are generated by the generator matrix of a $(\frac{M}{T}\alpha_i, \alpha_i)$ MDS code, so that they have full rank. Let the indices of the appeared variables be denoted by the vectors $\mathcal{I}_{x^{[k]}} \in \mathbb{N}^{TN^{K-1} \times 1}, \forall k \in [1:K]$. From Lemma 1, we have

$$x^{[k]}_{\mathcal{I}_{x^{[k]}}} \sim S_k[(1:TN^{K-1}), :]W_k. \tag{31}$$

Thus privacy is guaranteed.

Finally, we compute the ratio of the number of desired symbols to the number of total downloaded symbols (from N responding databases),

$$R = \frac{N}{N}\frac{\sum_{j=1}^{K}(N-T)^{j-1}T^{K-j}\binom{K-1}{j-1}}{\sum_{j=1}^{K}(N-T)^{j-1}T^{K-j}\binom{K}{j}} \tag{32}$$

$$= \frac{N}{N}\frac{N^{K-1}}{\frac{1}{N-T}\left[\sum_{j=1}^{K}(N-T)^{j}T^{K-j}\binom{K}{j}\right]} \tag{33}$$

[4]To be more specific, database symmetry refers to the property that each database downloads a equal number of instances for each type of sums, and message symmetry refers to the property that within each database, the symbols from each message are equivalent up to permutations. A more detailed treatment can be found in [25]. We initialize by downloading T^{K-1} symbols such that in Step 4 when we divide side information symbols, each database always obtains an integer number of side information symbols.

[5]The choice of α_i is to ensure both correctness and privacy. Specifically, it guarantees that over each layer, (1) sufficiently many undesired symbols are exposed to decode the remaining undesired symbols that interfere with the desired symbols, and (2) the number of symbols seen by any colluding set of databases matches the MDS code dimension such that they appear uniformly random. The proof appears later. For example, consider the setting in Section IV-B, where the desired message index $l = 1$. Consider the undesired message index $k = 2$. Here $\Delta = 2$, i.e., $\mathcal{K}_1 = \{2\}, \mathcal{K}_2 = \{2, 3\}$, $\alpha_1 = 12$ and $\alpha_2 = 6$.

[6]Over the j-th layer, the downloads are in the form of sums of j symbols, each from one distinct message. The term $(N-T)^{j-1}T^{K-j}$ comes from the side information exploitation step (Step 4) and can be verified recursively. A detailed analysis in similar flavor can be found in [25].

$$\begin{bmatrix} x^{[k]}_{\mathcal{K}_1} \\ x^{[k]}_{\mathcal{K}_1\cup\{l\}} \\ \hline x^{[k]}_{\mathcal{K}_2} \\ x^{[k]}_{\mathcal{K}_2\cup\{l\}} \\ \hline \vdots \\ \hline x^{[k]}_{\mathcal{K}_\Delta} \\ x^{[k]}_{\mathcal{K}_\Delta\cup\{l\}} \end{bmatrix} = \left[\begin{array}{cccc} \mathrm{MDS}_{\frac{M}{T}\alpha_1\times\alpha_1} & 0 & 0 & 0 \\ \hline 0 & \mathrm{MDS}_{\frac{M}{T}\alpha_2\times\alpha_2} & 0 & 0 \\ \hline 0 & \cdots & \ddots & 0 \\ \hline 0 & 0 & 0 & \mathrm{MDS}_{\frac{M}{T}\alpha_\Delta\times\alpha_\Delta} \end{array}\right] S_k[(1:TN^{K-1}),:]W_k$$

$$= \frac{\frac{1}{N}N^K}{\frac{1}{N-T}\left(N^K-T^K\right)} = \frac{1-\frac{T}{N}}{1-\frac{T^K}{N^K}} \tag{34}$$

$$= \left(1+\frac{T}{N}+\frac{T^2}{N^2}+\cdots+\frac{T^{K-1}}{N^{K-1}}\right)^{-1} \tag{35}$$

Thus, the PIR rate achieved by the scheme always matches the capacity.

Remark: When we set $T = 1$, $M = N$, Theorem 1 recovers the PIR capacity result in [25]. The two schemes achieve the same rate (capacity achieving), but the two differ in that although the query structures are the same, the specialization here uses MDS codes over a large field while the specialization in [25] uses permutations over message bits.

V. Proof of Theorem 1 and Theorem 2: Converse

Clearly the capacity of robust T-private PIR cannot be larger than the capacity of T-private PIR. Therefore, we only need to prove the converse for T-private PIR.

For compact notation, let us define

$$\mathcal{Q} \triangleq \{Q_n^{[k]} : k \in [1:K], n \in [1:N]\} \tag{36}$$

$$A_{\mathcal{I}}^{[k]} \triangleq \{A_n^{[k]} : n \in \mathcal{I}\} \tag{37}$$

$$\mathcal{H}_T \triangleq \frac{1}{\binom{N}{T}} \sum_{\mathcal{T}:|\mathcal{T}|=T} \frac{H(A_{\mathcal{T}}|\mathcal{Q})}{T}, \mathcal{T} \subset [1:N] \tag{38}$$

We first state Han's inequality ([35, Th. 17.6.1]), which will be used later and is described here for the sake of completeness.

Theorem 3 (Han's Inequality, [35, Th. 17.6.1]:)

$$\mathcal{H}_T \geq \frac{H(A_1^{[k]}, A_2^{[k]}, \cdots, A_N^{[k]}|\mathcal{Q})}{N} \tag{39}$$

We next proceed to the converse proof. The proof of outer bound for Theorem 1 is based on an induction argument. To set up the induction, we will prove the outer bound for $K = 1$ (the trivial case) for arbitrary N, T, and then proceed to the case of arbitrary K.

A. $K = 1$ Message, N Databases

$$L = H(W_1) = H(W_1|\mathcal{Q}) \tag{40}$$

$$= I(A_1^{[1]}, A_2^{[1]}, \cdots, A_N^{[1]}; W_1|\mathcal{Q}) \tag{41}$$

$$= H(A_1^{[1]}, A_2^{[1]}, \cdots, A_N^{[1]}|\mathcal{Q}) \tag{42}$$

$$\leq N\mathcal{H}_T \tag{43}$$

$$\leq \sum_{n=1}^{N} H(A_n|\mathcal{Q}) \tag{44}$$

$$\Rightarrow R = \frac{L}{D} \leq \frac{L}{\sum_{n=1}^{N} H(A_n|\mathcal{Q})} \leq 1 \tag{45}$$

where (43) follows from Han's inequality, and (44) is due to the property that dropping conditioning does not reduce entropy.

B. $K \geq 2$ Messages, N Databases

Consider $\mathcal{T} \subset [1:N]$ with cardinality $|\mathcal{T}| = T$. Denote the complement of $\mathcal{T}$ as $\overline{\mathcal{T}}$. From $A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[1]}, \cdots, A_{\overline{\mathcal{T}}}^{[K]}, \mathcal{Q}$, we can decode all K messages $W_1, \cdots, W_K$.

$$KL = H(W_1, \cdots, W_K|\mathcal{Q}) \tag{46}$$

$$= I(A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[1]}, \cdots, A_{\overline{\mathcal{T}}}^{[K]}; W_1, \cdots, W_K|\mathcal{Q}) \tag{47}$$

$$= H(A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[1]}, \cdots, A_{\overline{\mathcal{T}}}^{[K]}|\mathcal{Q}) \tag{48}$$

$$= H(A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[1]}|\mathcal{Q}) + H(A_{\overline{\mathcal{T}}}^{[2]}, \cdots, A_{\overline{\mathcal{T}}}^{[K]}|A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[1]}, \mathcal{Q}) \tag{49}$$

$$\leq N\mathcal{H}_T + H(A_{\overline{\mathcal{T}}}^{[2]}, \cdots, A_{\overline{\mathcal{T}}}^{[K]}|A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[1]}, W_1, \mathcal{Q}) \tag{50}$$

$$\leq N\mathcal{H}_T + H(A_{\overline{\mathcal{T}}}^{[2]}, \cdots, A_{\overline{\mathcal{T}}}^{[K]}|A_{\mathcal{T}}, W_1, \mathcal{Q}) \tag{51}$$

$$= N\mathcal{H}_T + H(A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[2]}, \cdots, A_{\overline{\mathcal{T}}}^{[K]}|W_1, \mathcal{Q}) - H(A_{\mathcal{T}}|W_1, \mathcal{Q}) \tag{52}$$

$$= N\mathcal{H}_T + H(A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[2]}, \cdots, A_{\overline{\mathcal{T}}}^{[K]}, W_2, \cdots, W_K|W_1, \mathcal{Q}) - H(A_{\mathcal{T}}|W_1, \mathcal{Q}) \tag{53}$$

$$= N\mathcal{H}_T + (K-1)L - H(A_{\mathcal{T}}|W_1, \mathcal{Q}) \tag{54}$$

where (50) is due to the fact that W_1 is a function of $(A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[1]}, \mathcal{Q})$. (53) follows from the fact that $W_2, \cdots, W_K$ is a function of $(A_{\mathcal{T}}, A_{\overline{\mathcal{T}}}^{[2]}, \cdots, A_{\overline{\mathcal{T}}}^{[K]}, \mathcal{Q})$. In (54), the second term is due to the fact that the answers are deterministic functions of the messages and queries, and the messages are independent.

Consider (54) for all subsets of $[1:N]$ that have exactly T elements and average over all such subsets. We have

$$N\mathcal{H}_T \geq L + \frac{1}{\binom{N}{T}} \sum_{\mathcal{T}:|\mathcal{T}|=T} H(A_{\mathcal{T}}|W_1, \mathcal{Q}) \tag{55}$$

To proceed, we note that for the last term of (55), conditioning on W_1, the setting reduces to a PIR problem with $K-1$

messages and N databases. Thus, (55) sets up an induction argument, which claims that for the K messages setting,

$$N\mathcal{H}_T \geq L\left(1+\frac{T}{N}+\cdots+\frac{T^{K-1}}{N^{K-1}}\right) \tag{56}$$

We have proved the basis cases of $K=1$ in (43). Suppose now the bound (56) holds for $K-1$. Then plugging in (55), we have that the bound (56) holds for K. Since both the basis and the inductive step have been performed, by mathematical induction, we have proved that (56) holds for all K. The desired outer bound follows as

$$\begin{aligned} R=\frac{L}{D} &\leq \frac{L}{\sum_{n=1}^{N} H(A_n|\mathcal{Q})} \leq \frac{L}{N\mathcal{H}_T} \\ &\leq \left(1+\frac{T}{N}+\cdots+\frac{T^{K-1}}{N^{K-1}}\right)^{-1} \end{aligned} \tag{57}$$

Thus, the proof of the outer bound is complete.

VI. Conclusion

We characterize the capacity of robust T-private PIR with arbitrary number of messages, arbitrary number of (responding) databases, and arbitrary privacy level. Let us conclude with a few observations. First, while in this paper we adopt the zero error framework, we note that our converse extends in a straightforward manner to the ϵ-error framework as well, where the probability of error is only required to approach zero as the message size approaches infinity. An outline of this extension is provided in the Appendix. Therefore, for robust T-private PIR, the ϵ-error capacity is the same as the zero error capacity. Second, recall that the capacity achieving scheme for PIR in our prior work [25] had a remarkable feature that if some of the messages were eliminated and the scheme projected onto a subset of messages, it remained capacity optimal for that subset of messages. The same phenomenon is observed for our achievable scheme for robust T-private PIR. On the other hand, an important point of distinction of the previous achievable scheme in [25] from the achievable scheme in this paper is that the former directly uses each available side information symbol individually, whereas here we need MDS coded side information (uncoded side information symbols do not suffice). This is because of the T-privacy constraint which simultaneously creates multiple perspectives of external side information depending upon which subset of databases decides to collude. Third, we note that in this paper we require perfect privacy (refer to (6), $I(Q_{\mathcal{T}}^{[k]}, A_{\mathcal{T}}^{[k]}, W_1, \cdots, W_K; k) = 0$). Similar to the ϵ-error relaxation, we may relax this to a σ-privacy constraint, where the information leaked about the desired message index vanishes as the message size grows. That is, we could replace the privacy constraint (6) by $I(Q_{\mathcal{T}}^{[k]}, A_{\mathcal{T}}^{[k]}, W_1, \cdots, W_K; k) \leq \sigma$, where σ approaches zero as the message size approaches infinity. It turns out that the capacity under σ-privacy is the same as the capacity under perfect privacy. Our converse proof extends to this setting by noting that the σ-privacy constraint implies that for any two message indices $k_1, k_2 \in [1:K]$, the difference $H(Q_{\mathcal{T}}^{[k_1]}, A_{\mathcal{T}}^{[k_1]}, W_1, \cdots, W_K) - H(Q_{\mathcal{T}}^{[k_2]}, A_{\mathcal{T}}^{[k_2]}, W_1, \cdots, W_K) = \sigma'$, vanishes with the message size and all other steps remain unchanged.

Finally, we note that since we focus only on download cost, upload cost is not optimized in this work. However, even with T-privacy, significant optimizations of upload cost are possible through refinements of our achievable scheme. For example, the symbols may be grouped in a manner that randomizations are needed only within smaller groups, which may reduce the number of possible queries, and the size of the field of operations significantly. To illustrate this, consider the achievable scheme for $K=2, N=3, T=2$ that was presented in Section IV-A, where each message is comprised of 9 symbols. We will operate over $\mathbb{F}_2$ (note that previously we required the field size q to be larger than 9). Suppose we divide the 9 bits into 3 groups of 3 bits each, and label the groups so that A_1 represents the first three bits of W_1, A_2 the next three and A_3 represents the last three bits from W_1. Similarly, let B_1, B_2, B_3 represent three groups of three bits each from W_2. Now, for any group of 3 bits, say $X=(x_1,x_2,x_3)$, let $X(1), X(2), X(3)$ represent three randomly chosen linearly independent elements from the set $\{x_1, x_2, x_3, x_1+x_2, x_1+x_3, x_2+x_3, x_1+x_2+x_3\}$, i.e., selected uniformly among the choices that do not sum to zero in $\mathbb{F}_2$. This essentially means that $X(1), X(2)$ may be freely chosen as any two distinct elements of the set and then $X(3)$ is chosen uniformly from the 4 elements that are not $X(1), X(2)$ or $X(1)+X(2)$. The queries are constructed as follows.

DB1	DB2	DB3
$A_1(1), A_2(1)$	$A_2(2), A_3(2)$	$A_3(3), A_1(3)$
$B_1(1), B_2(1)$	$B_2(2), B_3(2)$	$B_3(1+2), B_1(1+2)$
$A_3(1)+B_3(1)$	$A_1(2)+B_1(2)$	$A_2(3)+B_2(1+2)$

where we use the notation $X(1+2)=X(1)+X(2)$ for brevity. Note that for the undesired symbols B, we used the $(2,3)$ MDS code $(B(1), B(2)) \longrightarrow (B(1), B(2), B(1+2))$ within each group. Due to the grouping of symbols the upload cost is significantly reduced. Moreover, because of the grouping we are able to operate over a smaller field. Whereas the original scheme presented in Section IV-A uses $(6,9)$ MDS codes which do not exist over $\mathbb{F}_2$, the refined example presented above uses only a $(2,3)$ MDS code which does exist over $\mathbb{F}_2$. As illustrated by this example, optimizations of upload costs as well as symbol size remain interesting avenues for future work.

Appendix: ϵ-Error Capacity

In the ϵ-error framework, a rate R is said to be ϵ-error achievable if there exists a sequence of PIR schemes, indexed by the message size L, each of rate greater than or equal to R, for which the probability of error approaches 0 as $L \to \infty$. For such a sequence of PIR schemes, from Fano's inequality, we have the following correctness condition (corresponding to (5) in the zero-error framework).

$$\text{[Correctness]} \quad H(W_k|A_1^{[k]}, \cdots, A_N^{[k]}, Q_1^{[k]}, \cdots, Q_N^{[k]}) = o(L) \tag{58}$$

where any function of L, say $f(L)$, is said to be $o(L)$ if $\lim_{L\to\infty} f(L)/L = 0$. The supremum of ϵ-error achievable rates is called the ϵ-error capacity.

We next show that the zero error capacity results in both Theorem 1 and Theorem 2 hold under the ϵ-error framework.

Note that zero-error achievable schemes automatically satisfy ϵ-error criterion (58). Thus we are only left to prove that the converse proof extends to the ϵ-error setting. This follows from the simple observation that in the current zero-error converse proof, in (41) and (47), we can simply replace the zero-error correctness condition (5) with the ϵ-error correctness condition (58) and all other steps follow in the same manner.

REFERENCES

[1] Y. Gertner, Y. Ishai, E. Kushilevitz, and T. Malkin, "Protecting data privacy in private information retrieval schemes," in *Proc. 30th Annu. ACM Symp. Theory Comput.*, 1998, pp. 151–160.

[2] J. Feigenbaum, "Encrypting problem instances," in *Advances in Cryptology*. Berlin, Germany: Springer, 1985, pp. 477–488.

[3] M. Abadi, J. Feigenbaum, and J. Kilian, "On hiding information from an oracle," in *Proc. 19th Annu. ACM Symp. Theory Comput.*, 1987, pp. 195–203.

[4] D. Beaver and J. Feigenbaum, "Hiding instances in multioracle queries," in *Proc. STACS*, 1990, pp. 37–48.

[5] D. Beaver, J. Feigenbaum, J. Kilian, and P. Rogaway, "Locally random reductions: Improvements and applications," *J. Cryptol.*, vol. 10, no. 1, pp. 17–36, 1997.

[6] A. Shamir, "How to share a secret," *Commun. ACM*, vol. 22, no. 11, pp. 612–613, 1979.

[7] A. Beimel, Y. Ishai, E. Kushilevitz, and I. Orlov, "Share conversion and private information retrieval," in *Proc. 27th Annu. Conf. Comput. Complex.*, Jun. 2012, pp. 258–268.

[8] S. Yekhanin, "Locally decodable codes and private information retrieval schemes," Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, USA, 2007.

[9] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai, "Batch codes and their applications," in *Proc. 36th Annu. ACM Symp. Theory Comput.*, 2004, pp. 262–271.

[10] Y. Ishai and E. Kushilevitz, "On the hardness of information-theoretic multiparty computation," in *Advances in Cryptology*. Berlin, Germany: Springer-Verlag, 2004, pp. 439–455.

[11] H. Sun and S. A. Jafar. (2016). "Blind interference alignment for private information retrieval." [Online]. Available: https://arxiv.org/abs/1601.07885

[12] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, "On the locality of codeword symbols," *IEEE Trans. Inf. Theory*, vol. 58, no. 11, pp. 6925–6934, Nov. 2012.

[13] A. G. Dimakis, K. Ramchandran, Y. Wu, and C. Suh, "A survey on network codes for distributed storage," *Proc. IEEE*, vol. 99, no. 3, pp. 476–489, Mar. 2011. [Online]. Available: http://arxiv.org/abs/1004.4438

[14] Y. Birk and T. Kol, "Coding on demand by an informed source (ISCOD) for efficient broadcast of different supplemental data to caching clients," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2825–2830, Jun. 2006.

[15] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.

[16] W. Gasarch, "A survey on private information retrieval," *Bull. EATCS*, vol. 82, pp. 72–107, Feb. 2004.

[17] S. Yekhanin, "Private information retrieval," *Commun. ACM*, vol. 53, no. 4, pp. 68–73, 2010.

[18] R. Ostrovsky and W. E. Skeith, III, "A survey of single-database private information retrieval: Techniques and applications," in *Public Key Cryptography*. Berlin, Germany: Springer-Verlag, 2007, pp. 393–411.

[19] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, "Private information retrieval," in *Proc. 36th Annu. Symp. Found. Comput. Sci.*, Oct. 1995, pp. 41–50.

[20] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private information retrieval," *J. ACM*, vol. 45, no. 6, pp. 965–981, 1998.

[21] A. Ambainis, "Upper bound on the communication complexity of private information retrieval," in *Automata, Languages and Programming*. Berlin, Germany: Springer-Verlag, 1997, pp. 401–407.

[22] A. Beimel, Y. Ishai, E. Kushilevitz, and J.-F. Raymond, "Breaking the $O(n^{1/(2k-1)})$ barrier for information-theoretic private information retrieval," in *Proc. 43rd Annu. IEEE Symp. Found. Comput. Sci.* Nov. 2002, pp. 261–270.

[23] Z. Dvir and S. Gopi, "2-server PIR with sub-polynomial communication," in *Proc. 47th Annu. ACM Symp. Theory Comput.*, 2015, pp. 577–584.

[24] T. H. Chan, S.-W. Ho, and H. Yamamoto, "Private information retrieval for coded storage," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2015, pp. 2842–2846.

[25] H. Sun and S. A. Jafar, "The capacity of private information retrieval," *IEEE Trans. Inf. Theory*, vol. 63, no. 7, pp. 4075–4088, Jul. 2017.

[26] A. Beimel, Y. Ishai, and E. Kushilevitz, "General constructions for information-theoretic private information retrieval," *J. Comput. Syst. Sci.*, vol. 71, no. 2, pp. 213–247, 2005.

[27] O. Barkol, Y. Ishai, and E. Weinreb, "On locally decodable codes, self-correctable codes, and t-private PIR," *Algorithmica*, vol. 58, no. 4, pp. 831–859, 2010.

[28] A. Beimel and Y. Stahl, "Robust information-theoretic private information retrieval," *J. Cryptol.*, vol. 20, no. 3, pp. 295–321, 2007.

[29] Y. Gertner, S. Goldwasser, and T. Malkin, "A random server model for private information retrieval," in *Randomization and Approximation Techniques in Computer Science*. Berlin, Germany: Springer-Verlag, 1998, pp. 200–217.

[30] G. Fanti and K. Ramchandran, "Efficient private information retrieval over unsynchronized databases," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1229–1239, Oct. 2015.

[31] A. Beimel, Y. Ishai, and T. Malkin, "Reducing the servers computation in private information retrieval: PIR with preprocessing," in *Advances in Cryptology*. Berlin, Germany: Springer-Verlag, 2000, pp. 55–73.

[32] N. B. Shah, K. V. Rashmi, and K. Ramchandran, "One extra bit of download ensures perfectly private information retrieval," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun./Jul. 2014, pp. 856–860.

[33] A. Fazeli, A. Vardy, and E. Yaakobi, "Codes for distributed PIR with low storage overhead," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2015, pp. 2852–2856.

[34] R. Tajeddine, O. W. Gnilke, and S. E. Rouayheb. (2016). "Private Information Retrieval from MDS Coded Data in Distributed Storage Systems." [Online]. Available: https://arxiv.org/abs/1602.01458

[35] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 2006.

Hua Sun (S'12–M'17) received his B.E. in Communications Engineering from Beijing University of Posts and Telecommunications, Beijing, China, in 2011, M.S. in Electrical and Computer Engineering from University of California Irvine, USA, in 2013, and Ph.D. in Electrical Engineering from University of California Irvine, USA, in 2017. He is an Assistant Professor in the Department of Electrical Engineering at the University of North Texas, USA. His research interests include information theory and its applications to communications, privacy, networking, and storage.

Dr. Sun received the IEEE Jack Keil Wolf ISIT Student Paper Award in 2016, an IEEE GLOBECOM Best Paper Award in 2016, and the University of California Irvine CPCC Fellowship for the year 2011-2012.

Syed Ali Jafar (S'99–M'04–SM'09–F'14) received his B. Tech. from IIT Delhi, India, in 1997, M.S. from Caltech, USA, in 1999, and Ph.D. from Stanford, USA, in 2003, all in Electrical Engineering. His industry experience includes positions at Lucent Bell Labs, Qualcomm Inc. and Hughes Software Systems. He is a Professor in the Department of Electrical Engineering and Computer Science at the University of California Irvine, Irvine, CA USA. His research interests include multiuser information theory, wireless communications and network coding.

Dr. Jafar is a recipient of the New York Academy of Sciences Blavatnik National Laureate in Physical Sciences and Engineering, the NSF CAREER Award, the ONR Young Investigator Award, the UCI Academic Senate Distinguished Mid-Career Faculty Award for Research, the School of Engineering Mid-Career Excellence in Research Award, the School of Engineering Maseeh Outstanding Research Award, the IEEE Information Theory Society Best Paper Award, IEEE Communications Society Best Tutorial Paper Award, IEEE Communications Society Heinrich Hertz Award, and three IEEE GLOBECOM Best Paper Awards. His student co-authors received the IEEE Signal Processing Society Young Author Best Paper Award, and the Jack Wolf ISIT Best Student Paper Award. Dr. Jafar received the UC Irvine EECS Professor of the Year award six times, in 2006, 2009, 2011, 2012, 2014 and 2017 from the Engineering Students Council and the Teaching Excellence Award in 2012 from the School of Engineering. He was a University of Canterbury Erskine Fellow in 2010 and an IEEE Communications Society Distinguished Lecturer for 2013-2014. Dr. Jafar was recognized as a Thomson Reuters Highly Cited Researcher and included by Sciencewatch among The World's Most Influential Scientific Minds in 2014, 2015 and 2016. He served as Associate Editor for IEEE TRANSACTIONS ON COMMUNICATIONS 2004-2009, for IEEE COMMUNICATIONS LETTERS 2008-2009 and for IEEE TRANSACTIONS ON INFORMATION THEORY 2009-2012. He is a Fellow of the IEEE.